Beyond Spatial Pyramid Matching: Space-time Extended Descriptor for Action Recognition

Zhenzhong Lan, Alexander G. Hauptmann
Carnegie Mellon University
{lanzhzh, alex}@cs.cmu.edu

arXiv:1510.04565v1 [cs.CV] 15 Oct 2015

ABSTRACT

We address the problem of generating video features for action recognition. The spatial pyramid and its variants have been very popular feature models due to their success in balancing spatial location encoding and spatial invariance. Although it seems straightforward to extend spatial pyramid to the temporal domain (spatio-temporal pyramid), the large spatio-temporal diversity of unconstrained videos and the resulting significantly higher dimensional representations make it less appealing. This paper introduces the space-time extended descriptor, a simple but efficient alternative way to include the spatio-temporal location into the video features. Instead of only coding motion information and leaving the spatio-temporal location to be represented at the pooling stage, location information is used as part of the encoding step. This method is a much more effective and efficient location encoding method as compared to the fixed grid model because it avoids the danger of over committing to artificial boundaries and its dimension is relatively low. Experimental results on several benchmark datasets show that, despite its simplicity, this method achieves comparable or better results than spatio-temporal pyramid.

Categories and Subject Descriptors

H.4 [**Information Systems Applications**]: Miscellaneous

General Terms

Application

Keywords

Content-based video retrieval, visual feature, spatial-temporal pyramids, space-time extended descriptor

1. INTRODUCTION

With the constant expansion of on-line video collections, content-based video retrieval [9, 1, 26, 12] has become an important problem in multimedia retrieval. It is a difficult task since on-line videos are subject to large visual diversity. Robust to such variability, the Bag-of-Features (BoF) [23] model has been used as the main paradigm for representing videos. A BoF can be summarized as an encoding step and a pooling step [3]. Traditional pooling discards the local feature position information in the video space. However, this spatio-temporal information has been proven to be a discriminative cue [13]. Indeed, discriminative motion information is not equally distributed in the video space as shown in Figure 1. To benefit from this information, spatial pyramids [14, 13] divides a video using fixed grids and pools the features locally into each grid cell. Spatial pyramids can be easily extend to spatio-temporal pyramid (STP) to encode the order of actions or events, another important cue for video representation. For example, to perform an action called "sitting down", we need to gradually bend our knees and lower our body, while "standing up" is the reverse. Despite its usefulness, STP is not as effective as its spatial sibling on still images due to the fact that actions in unconstrained videos will have much more dramatic spatio-temporal variance than still images. For example, as illustrated in Figure 2, we see the person doing a cartwheel moves through out the frame whereas a still image would center the person in the frame. In this case, STP, which commits too much to the artificial boundaries may lead to a performance

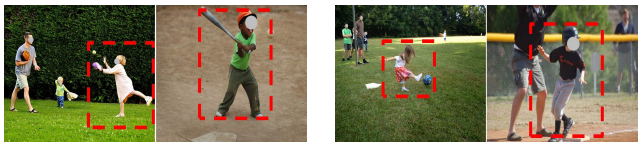

Figure 1: "Catch" and "hit" are likely to be distinguished by upper-bodies, especially hands while "kick" and "run" are more easily distinguished by legs.

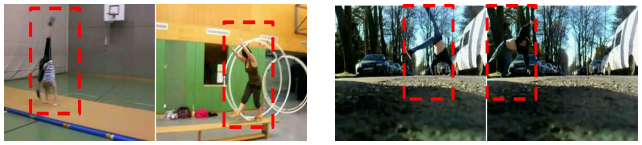

Inter-videos Intra-video

Figure 2: In different videos, actions localization can be subject to variation due to camera viewpoint change. But, even within a single video sequence, the action area can change among frames.

drop. Another major criticism of spatial pyramid is that it generates features with dimensions that are orders of magnitude higher than the original spatio-temporal invariant representations and hence make it computationally expensive to process. A more effective and efficient space-time encoding method is therefore critical for video representations in retrieving videos from large data collections.

In this work, we propose to take advantage of the spatio-temporal discriminative information with an emphasis on retaining the spatio-temporal robustness and controlling dimension explosion. Beyond standard spatial pooling which uses fixed segmentation grids, our proposed space-time extended descriptor (STED) augments the feature descriptor with spatio-temporal location information. By simultaneously encoding motion and location information, we remove the necessity of a complicated pooling stage and the danger of committing to artificial boundaries. Experimental results on several benchmark datasets show the advantages of STED over STP.

In the remainder of this paper, we start by providing more background information about spatial pyramids and its variants. We then compare STP and STED in detail. After that, an evaluation of our method is performed. Further discussions including potential improvements are given at the end.

2. RELATED WORK

Video retrieval research is conducted in a diverse setting where emphasis includes low-level and high-level feature design [26, 9], multi-modality fusion [1, 12] and multi-modality retrieval models [8]. Here we focus our review on low-level features design and encoding which is related to our latter experimental comparison. There has been a large amount of work in trying to build representations that keep spatial information of image patterns. Among them, spatial pyramid matching [14] is the most popular one. However, building spatial pyramids requires dimensions that are orders of magnitude higher than the original spatial invariant representations and hence make it less suitable for high dimensional encoding methods such as Fisher vector [20] and VLAD [2]. Spatial Fisher vector [10] and spatial augmentation [16, 22] provide more compact representations to encode spatial information and show similar performance as spatial pyramid methods. Few approaches consider encoding global temporal information into video representations. Oneata et al. [18] show that better action recognition performance can be achieved by dividing videos into two parts and encoding each one separately. Codella et al. [5] try to use temporal pyramid for event detection. They use n temporal segments, where n incrementally increases from 1 to 10.

3. SPACE-TIME ENCODING METHODS

In the following, we reformulate the spatial pyramid model and then extend this formulation to describe the STP and TED.

Let $D = \{d_1(\phi_1, x_1, y_1), d_2(\phi_2, x_2, y_2), ..., d_M(\phi_M, x_M, y_M)\}$ be a set of local feature descriptors extracted from a video. Each d_i contains the appearance/motion description ϕ_i, which is the typically local image descriptor, and the normalized pixel location (x_i, y_i) at which feature i is centered. We denote function $g : \phi_i \rightarrow \mathbb{R}^K$ as a local feature encoding scheme such as sparse-coding or locality encoding. Note that here g solely relies on the appearance/motion portion of the descriptors. We further denote by $G = \{G_1, ..., G_n\}$ a set of grid cells. Each G_j is a binary matrix indicating which pixel is active, $G_j \in 0, 1^{(s_x \times s_y)}$, (s_x, s_y) being the image size. Based on those definitions, we express the average spatial pooling operation as

$$X_j = \frac{1}{n} \sum_{i=1}^{M} G^j_{(x_i, y_i)} \times g(\phi_i). \quad (1)$$

3.1 Spatio-Temporal Pyramid (STP)

For temporal expansion, we simply add the temporal location into formula 1. That is

$$X_j = \frac{1}{n} \sum_{i=1}^{M} G^j_{(x_i, y_i, t_i)} \times g(\phi_i). \quad (2)$$

Note that by having one more location dimension, the resulting number of grid cells in G, hence number of dimension in feature X, will be orders of magnitude higher than the original spatial pyramid method. This may not be computationaly affordable for high dimensional representation such as Fisher Vector. For example, if we want to represent an Improved Dense Trajectory (IDT) [27] feature (dimension of ϕ_i is 426) using a Fisher vector representation with 256 Gaussian mixture models, then a $k \times k \times l$ with $k \in (1, 2)$ and $l \in (1, 3)$ STP can result in a representation about 4.4 million dimensions.

3.2 Space-Time Extended Descriptor (STED)

For STED, instead of using space-time information to locate the grid cells G, we use it to encode the feature. Formally, we have,

$$X = \frac{1}{n} \sum_{i=1}^{M} g(\phi_i, x_i, y_i, t_i). \quad (3)$$

One advantage of STED is that we avoid having to commit to artificial grid boundaries to define the spatial pooling regions, which can lead to very divergent representations for similar actions happenning in different space-time location. Another advantage is that the dimension of STED is only slightly higher than the original space-time invariant representation and is much lower than STP. Again taking the Dense Trajectory + Fisher vector setting for example, for each video, STED generates a feature with only about 0.21 million dimensions, which is about 20 times lower than the STP representation.

4. EXPERIMENTS

4.1 Experimental Setting

IDT with Fisher vector encoding [28] represents a current state-of-the-art for most real-world action recognition datasets. Therefore, we use it to evaluate our method. Note that although we use Fisher vector, our method can be applied to any quantization and pooling method such as VLAD [2]. Our baseline method uses the same settings as in [28]. These settings include the IDT feature extraction, Fisher vector representation and a linear SVM classifier.

	UCF50 (mAcc. %)		HMDB51 (mAcc. %)		Hollywood2 (mAP %)		Olympic (mAP %)	
l	Single	Pyramid	Single	Pyramid	Single	Pyramid	Single	Pyramid
1	**92.7**		59.6		65.8		**89.8**	
2	92.3	92.6	61.0	61.3	66.3	**67.4**	87.8	89.3
4	91.6	92.3	60.5	**61.6**	65.2	67.0	85.2	87.6
8	90.4	92.1	58.1	**61.6**	62.2	65.7	83.8	83.8

Table 1: Comparison of different temporal pyramid levels for STP.

IDT features are extracted using 15 frame tracking, camera motion stabilization with human masking and RootSIFT normalization and described by Trajectory, HOG, HOF and MBH descriptors. We use PCA to reduce the dimensionality of these descriptors by a factor of two. For Fisher vector representation, we map the raw feature descriptors into a Gaussian Mixture Model with 256 Gaussians trained from a set of randomly sampled 256000 data points. Power and L2 normalization are also used before concatenating different types of descriptors into a video based representation. For classification, we use a linear SVM classifier with a fixed C=100 as recommended by [28] and the one-versus-all approach is used for multi-class classification scenario.

For STP, we use $k \times k \times l$, k is 1 and 2 and $l \in (1, 2, 4, 8)$ grid cells. For STED, we attached the normalized 3 dimensional location to each descriptor as described in section 3.2.

4.2 Datasets

We use four action retrieval or classification datasets, UCF50, HMDB51, Hollywood2 and Olympic Sports, for evaluation. These datasets, which mainly involves actions, are selected because they are well known real-world datasets. The UCF50 dataset [21] has 50 action classes spanning over 6618 YouTube video clips that can be split into 25 groups. Video clips in the same group are generally very similar in background. Leave-one-group-out cross-validation as recommended by [21] is used and mean accuracy (mAcc) over all classes and all groups is reported. The HMDB51 dataset [11] has 51 action classes and 6766 video clips extracted from digitized movies and YouTube. [11] provides both original videos and stabilized ones. We only use original videos in this paper and standard splits with mAcc are used to evaluate the performance. The Hollywood2 dataset [15] contains 12 action classes and 1707 video clips that are collected from 69 different Hollywood movies. We use the standard splits with training and test videos provided by [15]. Mean average precision (mAP) is used to evaluate this dataset because multiple labels can be assigned to one video clip. The Olympic Sports dataset [17] consists of 16 athletes practicing sports, represented by a total of 783 video clips. We use the standard split with 649 training clips and 134 test clips and report mAP as in [17] for comparison purposes. Note that under this setting, on average, each class only has 8 testing samples, which may be too few to give a concrete measurement of the model.

4.3 Experimental Results

4.3.1 Spatial-Temporal Pyramids (STP)

In Table 1, we compare the performance of the STP at different temporal pyramid levels l. $l = 1$ corresponds to results only using a 1 and a 2×2 spatial pyramid pooling. From Table 1, we can see that, due to the large spatio-temporal diversity of unconstrained videos, the usefulness of STP is inconclusive. For some datasets such as HMDB51 and Hollywood2, it provides significant performance improvement while for other datasets such as UCF50 and Olympic Sports, it hurts the performance. Further division (from level 2 to 4 or 8) almost always results in worse performance with the exception of the pyramids in the HMDB51 dataset. These results show that a straightforward extension of spatial pyramids may hurt the performance.

STED	UCF50 (mAcc. %)	HMDB51 (mAcc. %)	Hollywood2 (mAP %)	Olympic (mAP %)
w/o	91.5	59.0	64.6	89.5
w	**93.0**	**62.1**	**67.0**	**89.8**

Table 2: Performance of STED.

4.3.2 Space-time Extension Descriptor (STED)

From Table 2, we see that, unlike STP, STED consistently improves the performance for all datasets. It is worth mentioning that both HMDB51 and Hollywood2 are very challenge datasets, more than 3% absolute improvement over space-time invariant representation is quite a notable gain. For Olympics dataset, because on average, each class only contains 8 testing examples, the improvement may not be statistically meaningful.

4.3.3 Comparing with the State-of-the-Art

In Table 3, we compare STP and STED, along with other recently published approaches. From Table 1, we see that a temporal level of 2 gives stable results for STP, therefore we set the temporal level as 2. From Table 3, we can see that STED achieves similar or better results in all four datasets. Considering that STED requires much lower dimensional encoding hence much less memory and computational cost than STP, it is a better space-time encoding method to use.

Note that although we list several recent approaches here for comparison purposes, they do use different features or settings. Shi et al. [24] use random sampled feature points and HOG, HOF, HOG3D and MBH descriptors. Jain et al. [7]'s approach incorporates a new motion descriptor. Oneata et al. [18] focus more on testing Spatial Fisher vector for multiple action and event tasks. Ciptadi et al. [4] presented a novel action representation based on encoding the global temporal movement of an action. The most comparable one is Wang et al. [28], from which we build our approaches and which serves as our baseline. Again, STED achieves comparable or better results than these state-of-art methods.

HMDB51 (MAcc. %)		Hollywood2 (MAP %)		UCF50 (MAcc. %)		Olympics Sports (MAP %)	
Oneata *et al.* [18]	54.8	Sapienz *et al.* [23]	59.6	Shi *et al.* [24]	83.3	Jain *et al.* [7]	83.2
Wang & Schmid [28]	57.2	Jain *et al.* [7]	62.5	Ciptadi *et al.* [4]	90.0	Adrien *et al.* [6]	85.5
Simonyan *et al.* [25]	57.9	Oneata *et al.* [18]	63.3	Oneata *et al.* [18]	90.0	Oneata *et al.* [18]	89.0
Peng *et al.* [19]	61.1	Wang & Schmid [28]	64.3	Wang & Schmid [28]	91.2	Wang & Schmid [28]	**91.1**
STP ($l = 2$)	61.3	STP($l = 2$)	**67.4**	STP ($l = 2$)	92.6	STP($l = 2$)	89.3
STED	**62.1**	STED	67.0	STED	**93.0**	STED	89.8

Table 3: Comparison of our results to the state-of-the-arts.

5. DISCUSSION

In this paper, we propose STED, a simple extension of local descriptors with spatio-temporal information. Despite it simplicity, STED is a much more effective and efficient way to to encode space-time location than previous methods. By simultaneously coding appearance, motion and location, STED avoids the danger of committing to artificial grid boundaries that define the spatial-temporal pooling regions and hence is better in dealing with unconstrained video that have large spatial-temporal motion diversity. We compare STED with STP, a straightforward way to extend spatial pyramids, and show that STED generates representations with much lower dimension while achieving similar or better results. Further improvements include determining optimal weighting for appearance/motion and location. Also, implementing STED with additional location encoding methods.

6. REFERENCES

[1] A. Amir, M. Berg, S.-F. Chang, W. Hsu, G. Iyengar, C.-Y. Lin, et al. Ibm research trecvid-2003 video retrieval system. *NIST TRECVID*, 2003.

[2] R. Arandjelovic and A. Zisserman. All about vlad. In *CVPR*, 2013.

[3] Y.-L. Boureau, F. Bach, Y. LeCun, and J. Ponce. Learning mid-level features for recognition. In *CVPR*, 2010.

[4] A. Ciptadi, M. S. Goodwin, and J. M. Rehg. Movement pattern histogram for action recognition and retrieval. In *ECCV*. 2014.

[5] N. C. Codella, A. Natsev, G. Hua, M. Hill, L. Cao, L. Gong, and J. R. Smith. Video event detection using temporal pyramids of visual semantics with kernel optimization and model subspace boosting. In *ICME*, 2012.

[6] A. Gaidon, Z. Harchaoui, and C. Schmid. Activity representation with motion hierarchies. *International Journal of Computer Vision*, 107(3):219–238, 2014.

[7] M. Jain, H. Jégou, and P. Bouthemy. Better exploiting motion for better action recognition. In *CVPR*, 2013.

[8] L. Jiang, T. Mitamura, S.-I. Yu, and A. G. Hauptmann. Zero-example event search using multimodal pseudo relevance feedback. In *ICMR*, 2014.

[9] Y.-G. Jiang, C.-W. Ngo, and J. Yang. Towards optimal bag-of-features for object categorization and semantic video retrieval. In *CIVR*, 2007.

[10] J. Krapac, J. Verbeek, and F. Jurie. Modeling spatial layout with fisher vectors for image categorization. In *ICCV*, 2011.

[11] H. Kuehne, H. Jhuang, E. Garrote, T. Poggio, and T. Serre. Hmdb: a large video database for human motion recognition. In *ICCV*, 2011.

[12] Z.-Z. Lan, L. Jiang, S.-I. Yu, S. Rawat, Y. Cai, et al. Cmu-informedia at trecvid 2013 multimedia event detection. In *NIST TRECVID*, 2013.

[13] I. Laptev, M. Marszalek, C. Schmid, and B. Rozenfeld. Learning realistic human actions from movies. In *CVPR*, 2008.

[14] S. Lazebnik, C. Schmid, and J. Ponce. Beyond bags of features: Spatial pyramid matching for recognizing natural scene categories. In *CVPR*, 2006.

[15] M. Marszalek, I. Laptev, and C. Schmid. Actions in context. In *CVPR*, 2009.

[16] S. McCann and D. G. Lowe. Spatially local coding for object recognition. In *ACCV*. 2013.

[17] J. C. Niebles, C.-W. Chen, and L. Fei-Fei. Modeling temporal structure of decomposable motion segments for activity classification. In *ECCV*. 2010.

[18] D. Oneata, J. Verbeek, C. Schmid, et al. Action and event recognition with fisher vectors on a compact feature set. In *ICCV*, 2013.

[19] X. Peng, L. Wang, X. Wang, and Y. Qiao. Bag of visual words and fusion methods for action recognition: Comprehensive study and good practice. *arXiv preprint arXiv:1405.4506*, 2014.

[20] F. Perronnin, J. Sánchez, and T. Mensink. Improving the fisher kernel for large-scale image classification. In *ECCV*. 2010.

[21] K. K. Reddy and M. Shah. Recognizing 50 human action categories of web videos. *Machine Vision and Applications*, 24(5):971–981, 2013.

[22] J. Sánchez, F. Perronnin, and T. De Campos. Modeling the spatial layout of images beyond spatial pyramids. *Pattern Recognition Letters*, 33(16):2216–2223, 2012.

[23] M. Sapienza, F. Cuzzolin, and P. H. Torr. Feature sampling and partitioning for visual vocabulary generation on large action classification datasets. *arXiv preprint arXiv:1405.7545*, 2014.

[24] F. Shi, E. Petriu, and R. Laganiere. Sampling strategies for real-time action recognition. In *CVPR*, 2013.

[25] K. Simonyan and A. Zisserman. Two-stream convolutional networks for action recognition in videos. *arXiv preprint arXiv:1406.2199*, 2014.

[26] C. G. Snoek and M. Worring. Concept-based video retrieval. *Foundations and Trends in Information Retrieval*, 2(4):215–322, 2008.

[27] H. Wang, A. Klaser, C. Schmid, and C.-L. Liu. Action recognition by dense trajectories. In *CVPR*, 2011.

[28] H. Wang, C. Schmid, et al. Action recognition with improved trajectories. In *ICCV*, 2013.